

THIS OPINION WAS NOT WRITTEN FOR PUBLICATION

The opinion in support of the decision being entered today (1) was not written for publication in a law journal and (2) is not binding precedent of the Board.

Paper No. 29

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Ex parte MARK YOUNG

Appeal No. 96-1432
Application 08/232,600¹

ON BRIEF

Before HAIRSTON, JERRY SMITH and CARMICHAEL, Administrative Patent Judges.

HAIRSTON, Administrative Patent Judge.

DECISION ON APPEAL

This is an appeal from the final rejection of claims 1 through 9.

The disclosed invention relates to reduction of cache flushing in a cache memory system. Using

¹ Application for patent filed April 25, 1994. According to appellant, the application is a continuation of Application 07/787,547, filed November 4, 1991, abandoned.

a write through with no write allocate approach, a cache memory controller in the cache

memory system reduces cache flushes on write cache misses by performing cache flushes only for cacheable physical memory locations.

Claim 7 is illustrative of the claimed invention, and it reads as follows:

7. An improved computer system comprising a central processing unit (CPU) coupled to a cache memory and a main memory executing a plurality of processes, wherein cache hits/misses are determined based on virtual addresses, some of said virtual addresses are alias addresses to each other, and cacheability of memory locations is determined as an integral part of virtual address to physical address translations, said improvement comprising an improved cache memory controller coupled to said CPU, said cache memory, and said main memory, that allows said CPU to update said cache and main memory with an improved write through with no write allocate approach that reduces cache flushes on write cache misses by waiting for the results of said cacheability determinations, the results of said cacheability determinations being available after the results of the corresponding cache hit/miss determinations, then without detecting for alias addresses of the virtual addresses, conditionally performing cache flushes for cache write misses only for cacheable memory locations.

The reference relied on by the examiner is:

Frink et al. (Frink), "A Virtual Cache-Based Workstation Architecture," 2nd IEEE Conference on Computer Workstations, Computer Society Press of the IEEE, 1988, pages 80 through 87.

Claims 1 through 9 stand rejected under 35 U.S.C. § 103 as being unpatentable over Frink.

Reference is made to the briefs and the answer for the respective positions of the appellant and the examiner.

OPINION

We have carefully considered the entire record before us, and we will reverse the obviousness rejection of claims 1 through 9.

The write policy normally used by the cache disclosed by Frink is write through, with write allocate (page 81, column 2). In a discussion of virtual address synonyms, Frink states (page 82, column 2) that “[v]irtual address synonyms, or aliases, exist when two or more virtual addresses point to the same physical data.” According to Frink (page 83, columns 1 and 2):

Another problem with virtual address synonyms exists if only a write-through *without* write-allocate policy is used. In this situation, Process 1 reads from address VA1, and the data is inserted in the cache at line L1. Process 2 then writes to address VA2 (which is a synonym to VA1 and indexes to line L1), but L1 is not affected because a write miss occurs as the result of a cache tag mismatch. When Process 1 reads VA1, it accesses stale data from line L1 instead of the new data that’s been written by Process 2.

To solve this problem, a write miss must do one of two things: invalidate the indexed cache entry, or update the entry to contain the new data.

According to the examiner (Answer, page 3), “‘A Virtual Cache-Based Workstation Architecture’ suggests that it was known to condition the flushing of a cache memory location on a write miss based on a cacheability determination at p. 82, c. 2, ll. 31-34 and p. 83, c. 2, ll. 7-11.”

Appellant argues (Brief, pages 7 through 9) that:

Although Virtual Cache discloses a virtually addressed cache, it does not teach or suggest conditioning the flushing of the affected cache line upon the determination of cacheability after a write miss. Virtual Cache merely teaches verifying cacheability

following a **read miss** rather than a write miss. Virtual Cache, p. 82 column 2 lines 4-17, and footnote. . . .Indeed, in a subsequent discussion of write misses, Virtual Cache **does not suggest checking the cacheability bit** of a physical memory location before invalidating a cache line. For example, after a write miss, Virtual Cache teaches that a cache line be **replaced** by the data that has just been written to memory, in effect flushing the previous cache line. Virtual Cache, p. 82 lines 31-34.

Even when it undertakes a discussion of synonymous virtual addressing (aliases), Virtual Cache teaches away from conditioning a cache line flush on a determination of cacheability.

We agree with appellant that Frink neither teaches nor would have suggested to one of ordinary skill in the art “conditioning a cache line flush on a determination of cacheability.” The examiner’s conclusion that Frink suggests such a teaching is completely without support in the record. For this reason, the obviousness rejection of claims 1 through 9 is reversed.

DECISION

The decision of the examiner rejecting claims 1 through 9 under 35 U.S.C. § 103 is reversed.

REVERSED

KENNETH W. HAIRSTON
Administrative Patent Judge

JERRY SMITH
Administrative Patent Judge

)
)
)
)
) BOARD OF PATENT
)
) APPEALS AND
)

Appeal No. 96-1432
Application 08/232,600

JAMES T. CARMICHAEL
Administrative Patent Judge

) INTERFERENCES
)
)
)

KWH/pgg
Jeffrey J. Blatt
Blakely, Sokoloff, Taylor & Zafman
12400 Wilshire BLVD., Seventh Floor
Los Angeles, CA 90025